

METHOD FOR SAN-BASED BOS INSTALL VOLUME GROUP

BACKGROUND OF THE INVENTION

1. Technical Field:

[0001] The present invention relates in general to volume groups and in particular to volume groups on storage area networks (SAN). Still more particularly, the present invention relates to providing efficient base operating system (BOS) installation for a volume group of a SAN.

2. Description of the Related Art:

[0002] Management of a computer system's physical memory, including distributed memory, involves dividing the physical memory into manageable sections, assigning physical addresses to those sections, and mapping those physical addresses to logical addresses. In some currently available computer systems, the management of the physical memory includes combining direct access storage devices (DASD) (also referred to as hard disks or physical volume (PV)) into groups called volume groups. Within each volume group, one or more logical volumes (LVs) are defined, which may be utilized for a number of system purposes, such as paging, storing raw data, or hosting a single file system.

[0003] Software known as logical volume managers (LVM), or simply volume managers, manages the fixed disk storage and the volume groups. The volume manager combines storage space on one or more hard disks into a single volume group and links the computer system's kernel to the volume group.

[0004] One example of a conventional LVM is the AIX LVM, which supports two volume group types: one with 32 disks and 256 logical volumes and one with 128 disks and 512 logical volumes, although further improvements are being made to accommodate larger numbers of disks and logical volumes. A volume manager such as AIX LVM provides virtualization between the physical disks and the user of the disk space such as a file system.

[0005] Each logical volume consists of one or more logical partitions (LPs), each corresponding to at least one physical partition, which is a fixed size region of space on a disk. If mirroring is specified for the logical volume, additional physical partitions are allocated to store the additional copies of each logical partition. These virtual logical disk devices appear to applications, databases, and file systems as if they were a physical disk device. However, a logical volume does not have the physical limitations of physical disk devices, and because of its virtual nature, a logical volume is not restricted to particular disks or a specific area of a disk. Data in logical volumes appears to be contiguous to the user but can be discontinuous on the physical volumes. This allows file systems, paging space, and other logical volumes to be re-sized or relocated, span multiple physical volumes, and have their contents replicated for greater flexibility and availability in the storage of data.

[0006] Because LVM is able to create logical volumes that appear contiguous in user space but which map to discontinuous spaces on disks, the disks are not required to be physically located within the system executing the LVM. In fact, the disks are often located remote from the computer system on which the LVM is executing and accessed via a network. The LVM then enables a system administrator to activate a particular logical volume for use by the client systems. For example, a file system is hosted on a logical volume, and once the boot process completes, the file system is available to be mounted by remote client systems via the network.

[0007] One network type that has combined the functionality of having multiple disk storage distributed across a network utilized for hosting file systems (and other applications) is a storage area network (SAN). Storage area networks are becoming more common in the computer network arena. A SAN comprises a network of storage disks that operate as a single resource and often links multiple servers together that each access the storage disks. An additional functionality of the SAN is that it enables data transfers between computers and disks at the same high peripheral channel speeds as when the two sets of components are coupled directly to each other. Thus, unlike conventional network file systems, which support client access to a file system over the Internet via transmission control protocol/Internet Protocol (TCP/IP), a SAN's file system appears as a directly connected storage device to the client, and the file system stored on a SAN is accessible without IP support. Also, although the storage disks may be located at

different locations, a SAN allows the system administrator to access the various disks from a single boot location.

[0008] Access to and operation of a file system and logical volume requires a boot device that activates the SAN. Further, the logical volumes each have a boot image as well as images of other files, etc. that are to be installed in a file system. The boot installation of devices within a SAN may be completed via several different techniques. In one conventional method, a CD ROM is provided with a boot image having a base operating system (BOS). This method is limited to a single medium (typically 700MB) or several similar media stacked together.

[0009] Currently, a boot device is identified in a nonvolatile random access memory (NVRAM), such as a CD or hard disk. Therefore, if a maintenance operation or installation requires the computer to boot from an alternative device, the identified device must be changed in NVRAM. For example, if an install operation is to be performed on a computer, an administrator may need to boot the computer from a removable medium, such as a compact disk (CD). Therefore, the administrator changes the boot disk in the computer to refer to a CD drive on the SAN.

[0010] In some conventional networks, a network install manager (NIM) is utilized to provide a network boot. With NIM, a first networked machine (server) requests the boot image from another server and installs individual pieces of software/products from the other server. Also, with NIM and with network file system (NFS), mounts installation/reading of the boot logical volume (bootlv) and installation of images in other networks occur at Internet Protocol (IP) speeds. This boot process requires the boot images and optional programming product OPP images be transmitted across the Internet in order to be installed and therefore incurred significant amounts of latency. Further, whenever the network is experiencing a down-time (i.e., network unavailable), NIM does not function and features such as disaster recovery cannot easily be accomplished.

[0011] Given the limitations of current implementations, which provide the boot image on a single removable CD that is also utilized to enable mounting of the file system, the present

invention recognizes that it would be desirable to create a volume group with a boot Device that is more easily accessible than via NIM installation across an IP network. The invention further realizes that it would be beneficial to provide the boot device within the SAN and remove reliance from external devices that may not be accessible when required. These and other benefits are provided by the invention described herein.

SUMMARY OF THE INVENTION

[0012] Disclosed is a method and system for providing reliable and quickly accessible backup boot/installation devices for volume groups and file systems on a storage area network (SAN). The invention provides a copy of the boot image and the base install images on multiple disks (physical volume) within a volume group of the SAN.

[0013] A boot logical volume "bootlv" is provided with a generic maintenance boot image mirrored on at least one disk (or set of disks) located out on the SAN. The disk(s) are encoded with similar boot functions as a conventional boot CD. The disk(s) provide physical storage for a volume group that has the base install images and all of the Optional Programming Product (OPP) images selected by the system administrator to be included in the install volume group. The volume group is able to boot up and run from any one of the disks since the boot image is mirrored/copied on each of the disks. Each disk has a particular identifier (ID) and SAN address by which it may be differentiated from the other disks on the SAN.

[0014] When the system administrator boots the computer system, the administrator points the system BIOS boot path at an install volume group disk with the boot image on it. The computer system boots into maintenance mode, and the operating system (OS) then initiates the install process. The install process imports the install volume group and installs the base operating system (BOS) from an image (i.e., bos.rte) within the volume group. The BOS image may be in one of several formats. In AIX, the format is that of a file in a file system in the volume group. The format may also be a raw logical volume. In these cases, the BOS image may just be overwritten with a new BOS image. Also, the boot image may be replaced by writing a new boot image to the bootlv. Then, the install process proceeds to install the proper device and OPP support.

[0015] In another embodiment, following system failure or where the administrator attempts system recovery from a corrupted boot image on the primary disk, the administrator points the BIOS boot path at a disk with the install volume group, initiates a boot-up in maintenance mode, and then imports the primary root volume group to rebuild the boot

image. A SAN-based volume group is thus allowed to quickly recover from failures, etc. without external software support.

[0016] Additionally, when the administrator desires to install new OPPs that is typically available from removable media, the administrator may import the install volume group, activate a mount of the file systems, and install the OPP image. Adding a new OPP to the volume group simply requires that those OPP image files be written to the mounted file system that contains the existing OPP images, and an update made to the table of contents file that exists for that file system. Following, the administrator un-mounts the file system and exports the volume group. Accordingly, base installation (bosinstall), failure recovery, and new OPP installation procedures are easily completed by importing a volume group and mounting file systems from within the volume group.

[0017] The above as well as additional objectives, features, and advantages of the present invention will become apparent in the following detailed written description.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] The novel features believed characteristic of the invention are set forth in the appended claims. The invention itself however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative embodiment when read in conjunction with the accompanying drawings, wherein:

[0019] **Figure 1** depicts a pictorial representation of a storage area network (SAN) data processing system in which the present invention may be implemented;

[0020] **Figure 2A** is a block diagram illustrating internal components of a data processing system within which several features of the invention may advantageously be implemented;

[0021] **Figure 2B** is a block diagram illustrating major functional components of a data processing system utilized as a SAN appliance in accordance with one embodiment of the invention;

[0022] **Figure 3** is a block diagram of a volume group subsystem that includes a SAN system, a plurality of physical disks, and client system to one illustrative embodiment of the invention;

[0023] **Figure 4** is a flow chart illustrating the process of selectively utilizing one of multiple BOS devices provided on a SAN according to one implementation of the invention;

[0024] **Figure 5** is a flow chart illustrating the process of installing new OPPs from removable media to storage devices of a SAN according to one implementation of the invention; and

[0025] **Figure 6** is a flow chart illustrating the process of completing recovery operations via an alternate boot device on the SAN according to one implementation of the invention.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENT(S)

[0026] The present invention provides a method and system for reliable and efficient boot up of a SAN-based volume group. by providing copies of base install images and a generic boot logical volume with a generic boot image on multiple disks (physical volumes) that make up the volume group. The boot image of the logical volume is copied and stored across the disks so that the disks look like a single device from a boot standpoint. During implementation, the installation process is made to point to a disk within the volume group, and the boot image for the appliance is copied from the disk within the volume group. The invention focuses on maintenance of volume groups and efficient (or fast) disaster/failure recovery. Three different implementations of the invention are provided, each corresponding to one of the flow charts described below.

[0027] Implementation of the invention does not limit the volume group to any particular size, and, in fact, the volume group is not even limited to a single medium. With the invention, a system administrator is able to initiate a boot process from any one of the disks within the volume group. The system administrator is also able to selectively switch from one disk to another among the multiple disks available within the volume group. This further allows the volume group to be activated without having to first mount the file system. From the perspective of the system being booted up, the activation process involves the system administrator identifying the volume group to the system and then loading the files from the selected disk(s) within the volume group to complete the boot process and install the files.

[0028] With reference now to the figures, and in particular **Figure 1**, which provides an illustration of a computer network comprising a storage area network (SAN) data processing system in which the present invention may be implemented. Computer network **100**, which may appropriately be referred to as SAN **100**, includes SAN data processing system (DPS) **105**, (which in one embodiment, may be a SAN appliance), and SAN fabric **102**, which is the medium used to provide communications links between various devices and computers connected together within SAN **100**. SAN fabric **102** may include connections, such as wire, wireless communication links, or fiber optic cables. In the depicted example, SAN fabric **102** may be a

fiber channel network, which is a high-speed transport technology utilized to build SANs. Of course, SAN 100 may also be implemented as a number of different types of networks, such as for example, an intranet, a local area network (LAN), or a wide area network (WAN). **Figure 1** is intended to serve solely as an example, and not as an architectural limitation for the present invention.

[0029] In the depicted example, SAN 100 includes machine A 104, machine B 106, and machine C 108, which may be servers and/or client workstations connected to SAN fabric 102. Also connected to SAN fabric 102 are a plurality of storage devices, including disk array 112, disk drives 114, tape drive 116, and optical drive 118. SAN 100 may include additional machines and other storage devices not shown. In some examples, thousands of such storage devices may be provided within SAN 100, particularly within the array 112 disks.

[0030] The invention is described from the perspective of a system manager performing boot up of SAN 100 and file systems stored on logical volumes of SAN 100. For simplicity of describing the invention, it is assumed that the server 120 is a SAN DPS 105 comprising a BOS boot image stored on at least one of the drives and utilized by the system administrator to complete the initial BOS install and other functional features of the invention. The invention is thus described with reference to the SAN DPS 105 (or simply SAN appliance) and boot operations controlled from SAN DPS 105. Specific configuration and operational features of a SAN appliance are described in the co-pending, related patent application Serial Number 10/448,238, entitled "Method, Apparatus, and Program for Performing Boot, Maintenance, or Install Operations on a Storage Area Network," filed on May 29, 2003. The entire content of that application is hereby incorporated by reference.

[0031] Client systems (e.g., 104-108) may boot from boot devices in the SAN 100. For example Machine A 104 may boot from a disk or other removable medium within disk drives 114 or tape drive 116 or optical drive 118. The address of the boot device is stored in non volatile random only memory (NVRAM) within machine A and, according to the invention, the specific boot device utilized for machine A 114 may be changed.

[0032] In accordance with a one embodiment of the present invention, SAN BOSBoot/BOSInstall device (SAN boot device) 120 associated with SAN DPS 105 is connected to SAN fabric 102. Each SAN DPS 105 has a built in function that allows the operating system (OS) to inquire of the SAN DPS 105 what boot device is to be utilized. Other functional features of the SAN DPS 105 when operating as a SAN appliance are described in detail in the related patent application, Serial Number 10/448,238, which has been previously incorporated herein by reference. Several of those functional features are described below.

[0033] A system administrator may set the boot type for a machine to normal boot, install, or maintenance boot. This boot type may be stored in SAN boot device 120. The administrator also stores lists of boot devices, rootvg devices, primary devices, secondary devices, and non-essential devices in the SAN boot device 120. The initial program load (IPL) firmware and network adapter firmware of client systems 104-108 are modified to query SAN boot device 120 for a list of devices. The address or world wide name (WWN) of the SAN boot device is stored within the network adapter of the client systems 104-108. When performing a basic operating system (BOS) boot operation or install operation, the firmware within the specific client system 104-108 sends a query for boot devices to the SAN boot device 120. Thereafter, the firmware may send similar queries for root volume group (rootvg) devices, primary devices, secondary devices, and nonessential devices. The SAN boot device 120 listens for queries from machines on the SAN boot 100. When a boot query is received, SAN boot device 120 returns an appropriate list of boot devices based on the boot type. Similarly, the SAN boot device 120 may also return lists of rootvg devices, primary devices, secondary devices, or nonessential devices based on the query received from machines on the SAN 100.

[0034] With reference now to **Figure 2A**, there is depicted a block diagram of a data processing system within which the present invention may be implemented. Data processing system 204 is an example of a client computer system and/or SAN DPS. Data processing system 204 employs a peripheral component interconnect (PCI) local bus architecture. Although the depicted example employs a PCI bus, other bus architectures such as Accelerated Graphics Port (AGP) and Industry Standard Architecture (ISA) may be used. Processor 202 and main memory 206 are connected to PCI local bus 207 through PCI bridge 208. PCI bridge 208 also may

include an integrated memory controller and cache memory for processor 202. Additional connections to PCI local bus 207 may be made through direct component interconnection or through add-in boards.

[0035] In the depicted example, SAN adapter 210, small computer system interface (SCSI) host bus adapter 212, and expansion bus interface 214 are connected to PCI local bus 206 by direct component connection. In contrast, graphics adapter 218, and audio/video adapter 219 are connected to PCI local bus 206 by add-in boards inserted into expansion slots. Expansion bus interface 214 provides a connection for a keyboard and mouse adapter 220, modem 222, and additional memory 224. SCSI host bus adapter 212 provides a connection for hard disk drive 226, tape drive 228, and CD-ROM drive 230. Typical PCI local bus implementations will support three or four PCI expansion slots or add-in connectors.

[0036] An operating system (OS) runs on processor 202 and is used to coordinate and provide control of various components within data processing system 204 in **Figure 2**. The operating system may be a commercially available operating system, such as Windows XP®, which is available from Microsoft Corporation. Instructions for the operating system, the object-oriented operating system, and applications or programs are located on storage devices, such as hard disk drive 226, and may be loaded into main memory 206 for execution by processor 202.

[0037] In an embodiment of the present invention, data processing system 200 boots from a device on a storage area network, completed via SAN adapter 210. Processor 202 loads boot code, also referred to as initial program load (IPL) firmware, from NVRAM 250. The boot code includes instructions for querying a SAN BOSboot/BOSinstall device for location of boot devices. When data processing system 200 is a client system, the boot/BOS install device is accessed via SAN adapter 210, which provides connection to a SAN DPS. When a SAN DPS BOS boot/install device is located within SAN adapter 210 or stored within memory 204. Firmware in SAN adapter 210 may include the world wide name (WWN) for the SAN storage device. Alternatively, the WWN for the SAN boot devices may be stored with the boot code in memory 204.

[0038] When performing a BOSboot operation or BOSinstall operation, the firmware sends an initial query for boot devices to the SAN boot device. Thereafter, the firmware may send similar queries for root volume group (rootvg) devices, primary devices, secondary devices, and non-essential devices. Thus, the present invention provides a dynamic mechanism for determining a set of devices to be configured for a given machine.

[0039] The depicted example in **Figure 2A** and above-described examples are not meant to imply architectural limitations on the invention. Those of ordinary skill in the art will appreciate that the hardware configuration in **Figure 2A** may vary depending on the implementation. For example, data processing system **204** also may be a notebook computer or hand held computer. Data processing system **204** also may be a kiosk or a Web appliance. Other internal hardware or peripheral devices, such as flash read-only memory (ROM), equivalent nonvolatile memory, or optical disk drives and the like, may be used in addition to or in place of the hardware depicted in **Figure 2A**.

[0040] **Figure 2B** is an exemplary block diagram of a SAN DPS with a SAN BOSboot/BOSinstall device in accordance with one embodiment of the present invention. SAN DPS **105** includes controller **252**, boot/install module **254**, communications (SAN) adapter **256**, and configuration module **308**. Communications adapter **316** may be a fibre channel adapter and is utilized for communicating with machines and devices on a SAN. SAN DPS **105** also stores boot device list **362**, root volume group (rootvg) device list **264**, primary device list **264**, and secondary device list **266**. More or fewer device lists may be stored within the scope of the present invention. For example, SAN DPS **250** may also store separate lists for normal boot devices, install boot devices, and maintenance boot devices.

[0041] The functional elements **252-258** and **262-268** are coupled to one another via the control/data signal bus **260**. Any architecture that facilitates the communication of control/data signals between elements **252-258** and **262-268** may be used without departing from the spirit and scope of the present invention.

[0042] Controller **252** controls the overall operation of the SAN DPS **105** and orchestrates

the operation of the other elements **254-258**. With the operation of the present invention, boot/install module **254** allows the system administrator to point to a boot device stored on a physical disk within the volume group. SAN DPS **105** may also store a boot type for each machine in boot device list **262** or elsewhere. When a query for boot devices is received, boot/install module identifies the appropriate boot devices in boot device list **262** based on the boot type for the requesting machine. When a query for rootvg devices is received, the boot/install module returns a list of rootvg devices from rootvg device list **264**.

[0043] Configuration module **258** instructs controller **252** to communicate with a machine to configure settings and to generate, update, delete, and modify device lists. For example, configuration module **258** may allow an administrator at a remote location to set a boot type for a specific machine or to add a device to the primary device list. In a preferred embodiment of present invention, configuration module **258** may include a Web server to provide a configuration interface through a Web browser at an administrator workstation.

[0044] The depicted example in Figure **2B** and above-description are not meant to imply architectural limitations. For example, SAN DPS **305** also may be a computer or server connected to the storage area network. The configuration module may also allow an administrator to set other options, such as the world-wide name (WWN) or address of the SAN DPS, a username and password for the SAN DPS, and an Internet-enabled small computer systems interface (Iscsi) adapter for connecting the disks over IP connection, etc.

[0045] The invention comprises a method and system by which a generic maintenance boot image is copied and stored on at least one disk (or set of disks) that is located out on the SAN, and then utilized for one of three applications. In one embodiment, the boot image is provided via conventional LVM mirroring. However, with other implementations, the boot image is provided by building a number of boot images for each disk. The disk(s) may be encoded with similar boot functions as a conventional boot CD and provide physical storage for a volume group that has the base install images and all of the OPP images selected by the system administrator to be included in the install volume group. In one embodiment, a volume group made up of several disks is able to boot up and run from any one of the several disks that

comprise the volume group since the bootlv is mirrored/copied on each of the disks. **Figure 3** provides more specific illustration of the functional components of a SAN DPS and interaction between the SAN DPS and the physical volumes, etc. to enable the functionality provided by the invention.

[0046] Turning now to **Figure 3**, there is illustrated a multi-disk volume group sub-system that includes mirrored boot devices according to one implementation of the invention. The volume group **307** is shown having two logical partitions, LP1 and LP2. LP1 hosts a file system, which is accessible by users and a system administrator via SAN connections to client systems **304** and SAN DPS **305**, respectively. Specifically, SAN fabric **302** provides the network connections for accessing the volume group **307**. Within SAN DPS **305** is illustrated the hardware and software components utilized to access the volume group **307**. These components include logical volume manager (LVM) **311**, OS **315** and file system controller **313**. Additionally, associated with OS **315** are boot device **320** and file images **321**, both of which are mirrored in physical drives **314** of the volume group **307** as described below. A system administrator's input/output (I/O) device **321** is provided to enable the system administrator to set up the file system and volume group, install software on the volume group, and when required, point the boot process away from the boot CD (disk drive) **320** to one of the physical volumes **314** available over the SAN fabric **302**. The I/O device **321** may be a client system with special permissions to enable system administrative updates or control of SAN DPS **305**.

[0047] The volume group **307** is made up of multiple physical volumes **314**, eight of which are shown labeled **PV1** through **PV8**. The actual number of physical volumes is not as important to an understanding of the invention; however, the invention requires at least one physical volume **314** within the volume group **307** contain a copy of the boot software code and images. The volume group **307** is a software abstraction of the combination of physical volumes **314**, which are combined and controlled by the LVM **311** to appear as a single congruent space. However, each disk retains its physical individuality and separate physical address, and the invention takes advantage of this fact by allowing the LVM **311** to (1) treat the physical disks as individual boot units over the SAN fabric **311** and (2) mirror boot code and images of the volume group **307** on the physical disks **314**. As indicated by the boot image storage **320B**

within each physical volume 314, copies of the boot code are built or mirrored/stored on each one of the physical volumes 314 of the volume group 307, which is accessible via SAN fabric 302. As described below, this configuration of the physical volumes 314 enables the boot code and OPP images to be read relatively quickly from any one of the disks at SAN speeds, and thus the features of the described invention apply to all SAN environments.

[0048] Once the boot process has completed, and the file system is up and running, the client 304 is able to mount the file system (FS) on the volume group 307 by communicating with the FS controller 313 via the SAN fabric 302. Although described with reference to specific embodiments and configuration, it is understood that the system of **Figure 3** is provided solely for illustrative purposes and not meant to be limiting on the invention.

[0049] The invention operates within the above hardware, firmware and software configuration to provide a method and system for providing more reliable backup boot/installation services for DPSs on a storage area network (SAN). As shown in **Figure 3**, the invention provides a boot image on every disk within a volume group of SAN. Copies of the base install images and a generic boot logical volume with a generic boot image are built or mirrored on multiple disks that make up the volume group. Once the copies have been made, the other processes of the invention may be implemented.

[0050] **Figure 4** is a flow chart illustrating the process of booting from one of the physical volumes of the volume group according to the invention. The process begins with the system administrator copying boot code to each (or some of the) disk that make up the volume group as shown at block 403. The system administrator accesses the SAN DPS and specifically the LVM and obtains a list of the physical volumes and their physical addresses. With this information, the system administrator selects one or more of the physical volumes/disks on which the boot code are to be mirrored or built. The administrator may choose to mirror/build the boot image on each of the physical disks; however, this may not be necessary when there are a large number of disks within the volume group, and only a few backup copies of the boot image are required. The copy of the boot image may be obtained from a physical CD-ROM boot disk or other removable type medium. Also, a special management GUI may be provided within the system

administrative tools that allow the system administrator to select the physical volumes of the volume group and initiate the mirroring of the boot image and OPP image. Notably, the mirroring occurs at SAN speeds and is thus completed relatively quickly.

[0051] Returning to **Figure 4**, the system administrator powers up the computer system as shown at block 405, and the computer goes into boot mode. That is, the computer system boots into maintenance mode, and when maintenance mode activates, the boot sequence generates a prompt for the system administrator to select a boot device. The system administrator selects a boot device on the volume group that is accessible via the SAN as shown at block 407. Each disk has a particular unique ID and SAN address by which it may be differentiated from the other disks on the SAN. The boot device is identified by a unique hardware identification (ID) utilized to select the device from among the list of physical volumes. The boot device selection feature is coded into the BIOS of the SAN DPS and, in one embodiment, provides the system administrator with a listing of install volume group disks with the boot image on them.

[0052] The computer system accesses the boot code on the selected physical volume via the SAN and initiates the boot and install processes from the physical volume as shown at block 409. The boot and install process imports the install volume group and installs the BOS image as indicated at block 411. The BOS image may be in one of several formats. In AIX, the format is that of a file in a file system in the volume group. The format may also be a raw logical volume. In these cases, the BOS image may just be overwritten with a new BOS image. Alternately, the boot image may be replaced by writing a new boot image to the bootlv. Then, the install process proceeds to install the proper device and OPP support. The system completes the boot and indicates system ready as shown at block 413.

[0053] In another embodiment, when the administrator wants to recover from a corrupted bootlv on the primary disk, the administrator simply directs the computer system at the correct disk with the install volume group, initiates a boot-up in maintenance mode, and then imports the primary root volume group to rebuild the boot image. **Figure 5** is a flow chart illustrating the process of completing disaster/failure recovery using the functionality of mirroring boot code on multiple disks with a volume group according to the invention. The mirroring of the boot code,

etc. is completed as provided in the above description of **Figure 4**. The failure recovery process is triggered responsive to the occurrence of a system failure as indicated at block **423**. The system is then powered back on in maintenance mode as shown at block **425**. The power on may be an automatic feature built into the OS or a manual feature initiated by a system administrator.

[0054] Following a system failure, the system may automatically point to the boot device on one of the install physical volumes across the SAN as indicated at block **427**. The boot code is initiated from the physical volume across the SAN and the image files are imported into the system as shown at block **429**. Then the images (device drivers, etc.) are installed on the computer system as shown at block **431**. Following the system boots up in normal mode and failure recovery process completes as shown at block **433**. The invention thus allows a SAN-based volume group to quickly recover from failures, etc. without external software support. In one embodiment, the disks are assigned a particular selection order in which they are accessed. During boot up, the first disk is accessed, and if the boot from the first disk fails, the second disk is automatically accessed, and so on.

[0055] **Figure 6** is a flow chart illustrating the process of installing new OPPs in a volume group (and/or DPS) using the functionality of mirroring boot code and images on multiple disks with a volume group according to one embodiment of the invention. The process begins when the system administrator desires to have new OPPs from removable media available for later install as shown at block **441**. The system administrator points the search device at the physical volume to import the install volume group as shown at block **443**. Then, the system administrator activates a mount of the file system as shown at block **444**, and imports the OPP from the physical volume as indicated at block **445**. This import of the OPP occurs via the SAN at SAN speed. Adding a new OPP to the volume group simply requires that those OPP image files be written to the mounted file system that contains the existing OPP images, and an update made of any table of contents files that exist for that file system. During installation, the installation code installs the products that is believed needed by the machine. The administrator may also select additional support to be installed prior to beginning the installation. After the BOS installation the administrator may import this volume group and install yet more products. The BOS provides support for automatic import, OPP install, and export of this volume group.

[0056] Returning to **Figure 6**, the image is installed and the OPP activated on the SAN DPS as shown at block 447. Once the OPPs are installed, the system administrator un-mounts the file system and exports the volume group as indicated at block 449. The administrator may select a menu option of install and specify the BOS install volume group as the source. The administrator would also provide a list of products to install. Following, the OS would import the volume group, mount the file system, install the products, dismount the file system, and export the volume group.

[0057] With the above three implementations, the invention provides base installation (BOSinstall), failure recovery, and new OPP installation procedures that now include importing a volume group and mounting file systems from the volume group. The advantage of the present invention is that the reading of the bootlv and installation of images occur at SAN speeds and not at internet protocol (IP) speeds as occurs with network install manager (NIM) or network file system (NFS) mounts. So boot images and OPP images are not being dragged across the internet to be installed on the DPS. Also, unlike with NIM systems which are rendered inoperable when the network is down, disaster recovery via the above SAN procedure is provided.

[0058] The invention further enables simultaneous use of a disk object by multiple different appliances connected to the volume group. This is possible because no single appliance has exclusive access to the group of disks making up the volume group and the boot image is mirrored on multiple disks within the volume group that may each be accessed by a different appliance. From a networked perspective, the invention provides a rotating set of disks that allows multiple systems to concurrently read the boot images. The volume group may also be opened as read only and take n reserve. The file system is mounted read only so that multiple machines may boot and install from the device concurrently. When utilized to provide new software support, the paths are updated and the updates are placed on the disk media before importing the new software packages.

[0059] One advantage of the present invention is that reading of the boot image and installation of images occur at SAN speeds and not at internet protocol (IP) speeds as occurs with network install manager (NIM) or network file system (NFS) mounts. Boot images and OPP

images do not have to be dragged across the Internet to be installed on the appliance. Also, unlike with NIM systems, which are rendered inoperable when the network is down, immediate disaster recovery via the above SAN procedure is available.

[0060] It is important to note that while the present invention has been described in the context of a fully functional data processing system, those skilled in the art will appreciate that the mechanism of the present invention is capable of being distributed in the form of a computer readable medium of instructions in a variety of forms, and that the present invention applies equally, regardless of the particular type of signal bearing media utilized to actually carry out the distribution. Examples of computer readable media include: nonvolatile, hard-coded type media such as Read Only Memories (ROMs) or Erasable, Electrically Programmable Read Only Memories (EEPROMs), recordable type media such as floppy disks, hard disk drives, CD-ROMs and DVD ROMs, and transmission type media such as digital and analog communication links, and wired or wireless communications links using transmission forms, such as, for example, radio frequency and light wave transmissions. The computer readable media may take the form of coded formats that are decoded for actual use in a particular data processing system.

[0061] While the invention has been particularly shown and described with reference to a preferred embodiment, it will be understood by those skilled in the art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention.